Lifted Relational Algebra with Recursion and Connections to Modal Logic

Eugenia Ternovska

arXiv:1612.09251v1 [cs.LO] 29 Dec 2016

ABSTRACT

We propose a new formalism for specifying and reasoning about problems that involve heterogeneous "pieces of information" – large collections of data, decision procedures of any kind and complexity and connections between them. The essence of our proposal is to lift Codd's relational algebra from operations on relational tables to operations on classes of structures (with recursion), and to add a direction of information propagation. We observe the presence of information propagation in several formalisms for efficient reasoning and use it to express unary negation and operations used in graph databases. We carefully analyze several reasoning tasks and establish a precise connection between a generalized query evaluation and temporal logic model checking. Our development allows us to reveal a general correspondence between classical and modal logics and may shed a new light on the good computational properties of modal logics and related formalisms.

1. INTRODUCTION

Our goal is to develop a formalism for linking various pieces of information, including powerful solvers. We take our inspiration in one the best-known and most influential formalisms in the Database community.

In 1970 Edgar (Ted) F. Codd introduced a relational data model and two query languages: relational calculus and relational algebra. Relational calculus is what we usually call first-order logic. The key contribution of Codd was to associate with a declarative specification language (first-order logic), a procedural counterpart which is the relational algebra, that later was implemented by smart engineers and became a multi-billion dollar industry of relational database management systems (RDBMS).

A significant change has happened in the past decade. While, at the low level, everything data-related boils down to SQL queries, *interactions between "larger", combinatorially harder, pieces became increasingly important.* Such "larger pieces" are business enterprises, knowledge bases, web services, software, solvers in the world of declarative problem solving, collections of learned data, potentially with numeric ranking assigned, etc. In particular, combinations of powerful solvers is already common in several communities, but has yet to make its way into the Database world.

While database queries, expressed using Codd's relational calculus, can be viewed as *relations definable with respect to a structure* (a database), declarative problem specifications can be understood as *axiomatizations of classes of structures.* The two notions (in italic) are defined in two consecutive chapters in the classic textbook on mathematical logic [4].

Our **first main idea** is to lift Codd's algebra from operations on relational tables to operations on *classes of structures* (which we call modules), and to add recursion. Each atomic module can be given, e.g., by a set of constraints in a constraint formalism that has an associated solver, such as Answer Set Programming, or Constraint Satisfaction Problem, or Integer Linear Programming. It can even be an agent or a business enterprise, a collection of databases, or a database that is only partially visible from the outside. Regardless of how modules are specified, our lifted relational algebra views them as atomic entities.

Binary relations, or relations with some kind of directionality, are very common. Problem solving often involves finding solutions for given inputs. Most combinatorial problems are of that form, many software programs and hardware devices are of that form. Graph databases use binary relations to link data. Role specifications in Description Logics, which are used to describe ontologies, some constructs in Datalog$^+_-$, temporal connections and specifications of action effects have directionality or describe information propagation. Our **second main idea** is to add information flow to the algebra. By doing so, we force expressions with multiple variables into binary and produce a calculus of binary relations. Modules with information propagation become (non-deterministic) actions. Modules without information propagation become propositions. Our algebra with information flow is exactly like our first algebra, that is, like Codd's relational algebra (with recursion), but has information propagation added.

The combination of the two ideas allows us to demonstrate a general connection between classical and modal logics through investigating the effects of adding information flow. We show that many formalisms for efficient reasoning are instances of the same phenomenon, which is also responsible for good computational properties of modal logics. Through our detailed study of several reasoning tasks, we provide an explanation of robust decidability of modal logics.

Structure of the Paper

In Section 2, we define our lifted version of Relational Algebra with recursion. We call this version of the algebra "flat" to distinguish it from the version with information propagation. We use a model-theoretic semantics, although various generalizations are possible. In Section 3, we define several reasoning tasks, in particular, Model Expansion that is responsible for adding information propagation.

In Section 4, we define our algebra with information flow, that we also call Dynamic Algebra or a Calculus of Binary Relations. In this algebra, second-order variables provide an easy way to model data flow. We interpret the Dynamic Algebra over transition systems where states are structures. We show that many interesting operations such as unary negation, constant tests and subexpression tests are definable in the algebra. We also discuss sequential composition and the reverse operations. Using the Dynamic Algebra, we produce a modal temporal logic $L\mu\mu$ and prove the equivalence of this modal logic to the calculus of binary relations. We show that several formalisms, an expressive Description Logic, Dynamic Logic, and Datalog$^{\pm}_{\neg}$ also use information propagation and can be viewed as fragments of $L\mu\mu$.

In Section 5, that we call *Queries, Machines, Modalities*, we introduce a model of computation (simple abstract machines) that are suitable for declarative problem solving. The machines are similar to Abstract State Machines of Yury Gurevich, however, in our machines, not only states, but also actions are structures. We discuss several important complexity measures. Formulae of our modal logic are programs for these machines. We describe reaching a solution to an algebraic specification in terms of these new devices. We then define a generalization of the Query Evaluation problem. Our main result is an equivalence between the reachability in the execution graph, the general Evaluation problem and and temporal model checking. The property holds under any assignment of the direction of information propagation to the internal relational variables. We then provide our explanation of why modal logics are so robustly decidable.

2. "FLAT" ALGEBRA

We call this version of the algebra "flat" to distinguish it from the version with information propagation below.

2.1 Syntax of the "Flat" Algebra

Let $\tau_M = \{M_1, M_2, \dots\}$ be a fixed vocabulary of *atomic module symbols*. Atomic module symbols are of the form $M_i(X_{i_1}, \dots, X_{i_k})$, (also written $M_i(\bar{X})$), where each X_i is a relational variable. Each X_j has an associated arity a_j. The set $\{X_{i_1}, \dots, X_{i_k}\}$ is called the *variable vocabulary of* M_i and is denoted $vvoc(M_i)$. Modules in τ_M are *atomic*. Modules that are not atomic are called *compound*. Algebraic expressions for *modules* are built by the grammar:

$$E ::= \bot | M_i | Z_j | E \cup E | -E | \pi_\delta E | \sigma_\Theta E | \mu Z_j.E. \quad (1)$$

Here, Z_j is a *module variable*. It must occur positively in the expression E, i.e., under an even number of the complementation $(-)$ operator. The set-theoretic operations are union $(\cup)$ and complementation $(-)$. Intersection $(\cap)$ and set difference are expressible. Projection $(\pi_\delta E)$ is a family of unary operations, one for each possible set of relational variables δ. Each symbol in δ must appear in E. The condition Θ in selection $\sigma_\Theta E$ is an expression of the form $L_1 \equiv L_2$, where L_i is a relational variable or 'R', where R is a relation (set of tuples of domain elements).[1] Thus, we bring semantic elements into syntax, and they became constants in the language. The operations (except $\mu Z_j.E$) are essentially like in Codd's relational algebra, except we include full complementation $(-)$ instead of set difference for generality. However, the operations are on objects of a higher order – on *classes of structures* rather than on relational tables.

REMARK 1. *While we follow an algebraic approach in presenting the syntax, it will be seen from the semantics that* **the constructs in this paper work the same way as the corresponding constructs in logic.**

The formalism is equivalent to a "lifted" version of first-order logic with the least fixed point operator (FO(LFP)). The "lifting" means that instead of regular predicate symbols, we have modules who's computational power we can control, and instead of object variables that range over domain elements, we have second-order variables ranging over relations. In particular, projection is onto a set of *relational* variables, selection is equality of relations rather than of domain elements. Note that first-order quantification is "encapsulated" in an axiomazation of each atomic module (e.g., an ILP or an ASP module) and is not visible from the outside of that module. Another way of viewing our formalism is SO+LFP−FO, that is, second-order logic with least fixed point "without" the first-order part.

2.2 Examples

We now give several examples of combinations of modules (pieces of information) in the algebra. For simplicity, we use common combinatorial problems since they are very familiar to most readers. A combinatorial problem can be viewed as a class of structures.[2]

EXAMPLE 1. Let $M_{\mathrm{HC}}(V, X, Y)$ and $M_{2\mathrm{Col}}(V, X, Z, T)$ be atomic modules "computing" a Hamiltonian Circuit and a 2-Colourability. They can do it in different ways. For example, M_{HC} can be an Answer Set Programming program, and $M_{2\mathrm{Col}}$ be an imperative program or a human child with two pencils. Here, V is a relational variable of arity 1, X, Y are relational variables of arity 2, and the first module decides if Y forms a Hamiltonian Circuit (represented as a set of edges) in the graph

[1] A more general version allows Θ to be built up using Boolean operations from the equivalence and non-equivalence operators, $\equiv, \not\equiv$. That choice of Θ may be more convenient to implement, but does not add expressive power since it is expressible trough the other operations.

[2] A module is a *set* of structures when the domain is fixed.

given by vertex set V and edge set X. Variable X of the second module has arity 2, and variables Z, T are unary; the module decides if unary relations Z, T specify a proper 2-colouring of the graph with edge set X. The following algebraic expression determines a combination of 2-Colouring and Hamiltonian Circuit, that is whether or not there is a 2-colourable Hamiltonian Circuit.[3]

$$M_{2\text{Col}-\text{HC}}(V, X, Z, T) := \pi_{V,X,Z,T}((M_{\text{HC}}(V, X, Y) \cap M_{2\text{Col}}(V, Y, Z, T)). \quad (2)$$

Projection "keeps" V, X, Z, T and hides the interpretation of Y in M_{HC}, since it is the same as Y's in $M_{2\text{Col}}$.

Example 2. This modular system can be used by a company that provides logistics services (arguments of atomic modules are omitted).

$$M_{LSP} := \sigma_{B \equiv B'}(M_K \cap M_{TSP}).$$

It decides how to pack goods and deliver them. It solves two NP-complete tasks interactively, – Multiple Knapsack (module M_K) and Travelling Salesman Problem (module M_{TSP}). The system takes orders from customers (items to deliver, their profits, weights), and the capacity of trucks, decides how to pack items in trucks, and for each truck, solves a TSP problem. The feedback B' about solvability of TSP is sent back to M_K. The two sub-problems, M_K and M_{TSP}, are solved by different sub-divisions of the company (potentially, with their own business secrets) that cooperate towards the common goal. A solution to the compound module, M_{LSP}, to be acceptable, must satisfy both sub-systems.

In some specifications, the use of a **recursive construct** is essential. For example, we may need to specify a recursive algorithm, or the semantics of the satisfaction relation in a logic, which is given by an inductive definition.

Example 3. In this example, we need recursion to specify a Dynamic Programming algorithm on a tree decomposition of a graph. The modules we use are M_{TD} that performs tree decomposition of each graph, and $M_{3\text{Col}}$ that is used recursively to performs 3-Colouring on each bag of the decomposition. The problem is represented as

$$M_{\text{TD}} \cap \mu Z.\Psi(Z, M_{\text{TD}}, M_{3\text{Col}}),$$

where Z is a module variable over which recursive iteration is performed, the least fixed point expression $\mu Z.\Psi(Z, M_{\text{TD}}, M_{3\text{Col}})$ specifies the dynamic programming algorithm. Details are omitted because of lack of space.

2.3 Valuations

To interpret algebraic expressions, we use valuations. This is an important notion used throughout the paper.

[3]We use := for "is by definition".

Definition 1 (Valuation). *Valuation $(v, \mathcal{V})$ is a pair of functions. Function v maps relational variables in $vvoc(M_i)$ to symbols in a relational vocabulary τ so that the arities of the relational variables in $vvoc(M_i)$ match those of the corresponding symbols in τ. Function $\mathcal{V}$ is parameterized by v and provides a domain (which does not have to be finite), and interpretations of atomic modules M_i as follows. Let V be the set of all τ-structures over the domain fixed by $\mathcal{V}$. Valuation $\mathcal{V}$ maps each atomic module symbol M_i to a subset $\mathcal{V}(v, M_i)$ of V so that for any two τ-structures $\mathcal{A}_1$, $\mathcal{A}_2$ which coincide on $vocab(M_i)$, we have $\mathcal{A}_1 \in \mathcal{V}(v, M_i)$ iff $\mathcal{A}_2 \in \mathcal{V}(v, M_i)$.*

All of the above applies to module variables Z_j as well. In practice, $\mathcal{V}$ can, for example, associate one module symbols with stable models of an ASP program, another module symbol with models of an ILP encoding, yet another one with a set of databases used by a particular enterprise, etc. We will also see, in part 4.6.2, that it is also possible to treat each modules simply as a predicate symbol.

Remark 2. *Valuations $\mathcal{V}$ (parameterized with v) can be viewed as "oracles" or decision procedures associated with modules, and can be of arbitrary computational complexity.*

2.4 Semantics of the "Flat" Algebra

The *extensions* $[\![E]\!]^{\mathcal{V},v}$ of algebraic expressions E are subsets of V (the set of all τ-structures over the domain fixed by $\mathcal{V}$) defined as follows.

$[\![\bot]\!]^{\mathcal{V},v} := \varnothing$.
$[\![M_i]\!]^{\mathcal{V},v} := \mathcal{V}(v, M_i)$ for some v.
$[\![Z_j]\!]^{\mathcal{V},v} := \mathcal{V}(v, Z_j)$ for some v.
$[\![E_1 \cup E_2]\!]^{\mathcal{V},v} := [\![E_1]\!]^{\mathcal{V},v} \cup [\![E_2]\!]^{\mathcal{V},v}$.
$[\![-E]\!]^{\mathcal{V},v} := V \backslash [\![E]\!]^{\mathcal{V},v}$.
$[\![\pi_\delta(E)]\!]^{\mathcal{V},v} := \{\mathcal{A} \mid \exists \mathcal{A}' \ (\mathcal{A}' \in [\![E]\!]^{\mathcal{V},v} \text{ and } \mathcal{A}|_\delta = \mathcal{A}'|_\delta)\}$.
$[\![\sigma_{L_1 \equiv L_2} E]\!]^{\mathcal{V},v} := \{\mathcal{A} \mid [\![E]\!]^{\mathcal{V},v} \text{ and } L_1^{\mathcal{A}} = L_2^{\mathcal{A}}\}$.
$[\![\mu Z_j.E]\!]^{\mathcal{V},v} := \bigcap \{\mathcal{E} \subseteq \mathcal{C} \mid [\![E]\!]^{\mathcal{V}[Z:=\mathcal{E}],v} \subseteq \mathcal{E}\}$.

Here, $\mathcal{V}[Z{:=}\mathcal{E}]$ means a valuation that is exactly like $\mathcal{V}$ except Z is interpreted as $\mathcal{E}$. Note that projection restricts each structure $\mathcal{B}$ of M to $\mathcal{B}|_\delta$ leaving the interpretation of other symbols open. Thus, it *increases* the number of models. Selection *reduces* the number of models.
This algebra may look very different from Codd's algebra because all modules are sets of τ-structures, that is, sets of tuples of the same length, while in Codd's algebra the length of tuples varies. There is no contradiction here – Codd's algebra is just what is seen through the "window of variables".

Definition 2 (Satisfaction, "flat" algebra). *Given a well-formed algebraic expression E defined by (1), we say that structure $\mathcal{A}$* satisfies *E under valuation $(\mathcal{V}, v)$, notation*

$$\mathcal{A} \models_{(\mathcal{V},v)} E,$$

if $\mathcal{A} \in [\![E]\!]^{\mathcal{V},v}$.

REMARK 3. *Note that while individual modules are already capable of solving optimization tasks (the optimum value can be given as an output in one of the arguments), the least fixed point construct can generate the least value over a collection of modules combined in an algebraic expression.*

2.5 Representation in Logic

PROPOSITION 1 (LOGIC COUNTERPART). *The algebraic operations are equivalently representable in logic, where '$\cup$' corresponds to disjunction, '$-$' to negation, 'π_ν' to second-order existential quantification over $\tau\backslash\nu$, 'σ_Θ' to conjunction with Θ, $\mu Z.E$ to the least fixed point construct.*

EXAMPLE 4. Expression (2) for $M_{2\text{Col}-\text{HC}}(V, X, Z, T)$ is represented in logic as

$$\exists Y[M_{\text{HC}}(V, X, Y) \wedge M_{2\text{Col}}(V, Y, Z, T)].$$

Note that first-order variables can be mimicked with second-order variables over singleton sets.

PROPOSITION 2. *When all relations are unary and the sets that interpret them are are singletons, the formalism collapses to FO(LFP).*

3. MODEL EXPANSION, RELATED TASKS

Model expansion [9] is the task of expanding a structure to satisfy a specification (a formula in some logic).

3.1 Three Related Tasks: Definitions

For a formula ϕ in any logic $\mathcal{L}$ with model-theoretic semantics, we can associate the following three tasks (all three for the same formula), satisfiability (SAT), model checking (MC) and model expansion (MX). We now define them for the case where ϕ has no free object variables.

DEFINITION 3 (Satisfiability (SAT_ϕ)). *Given: Formula ϕ. Find: structure $\mathcal{B}$ such that $\mathcal{B} \models \phi$. (The decision version is: Decide: $\exists\mathcal{B}$ s.t. $\mathcal{B} \models \phi$?)*

DEFINITION 4 (Model Checking (MC)). *Given: Formula ϕ, structure $\mathcal{A}$ for $vocab(\phi)$. Decide: $\mathcal{A} \models \phi$?*

There is no search counterpart for this task.
The following task (introduced in [9]) is at the core of this paper. The decision version of it is of the form "guess and check", where the "check" part is the model checking task we just defined.

DEFINITION 5 (Model Expansion (MX^σ_ϕ)). *Given: Formula ϕ with designated input vocabulary $\sigma \subseteq vocab(\phi)$ and σ-structure $\mathcal{A}$. Find: structure $\mathcal{B}$ such that $\mathcal{B} \models \phi$ and expands σ-structure $\mathcal{A}$ to $vocab(\phi)$. (The decision version is: Decide: $\exists\mathcal{B}$ such that $\mathcal{B} \models \phi$ and expands σ-structure $\mathcal{A}$ to $vocab(\phi)$?)*

Vocabulary σ can be empty, in which case the input structure $\mathcal{A}$ consists of a domain only. When $\sigma = vocab(\phi)$, model expansion collapses to model checking, $\mathsf{MX}^\sigma_\phi = \mathsf{MC}_\phi$. Note that, in general, the domain of the input structure in MC and MX can be infinite.
Let ϕ be a sentence, i.e., has no free object variables. Data complexity [12] is measured in terms of the size of the finite active domain. Data complexity of MX lies in-between model checking (full structure is given) and satisfiability (no structure is given).

$$\mathsf{MC}_\phi \leqslant \mathsf{MX}^\sigma_\phi \leqslant \mathsf{SAT}_\phi.$$

Of course, we consider the decision versions of the problems here. For example, for FO logic, MC is non-uniform AC^0, MX captures NP (Fagin's theorem), and SAT is undecidable. In SAT, the domain is not given. In MC and MX, the (active) domain is always given, which significantly reduces the complexity of these tasks compared to SAT. The relative complexity of the three tasks for several logics has been studied in [8].

3.2 Similar Task: Query Evaluation

In database literature, a task similar to MXis query evaluation. We define it for non-boolean queries, i.e., those with free object variables $\bar{y}$.

DEFINITION 6 (Query Evaluation (QE)). *Given: Formula $\phi(\bar{x})$ with free object variables $\bar{x}$, tuple of domain elements $\bar{a}$ and σ-structure $\mathcal{A}$, where $\sigma = vocab(\phi)$. Decide: $\mathcal{A} \models \phi[\bar{a}/\bar{x}]$?*

The task specifies the combined complexity of query evaluation. Data complexity requires formula ϕ to be fixed. To analyze expression complexity, we fix the database [12].
For *the same logic*, these two tasks, MX and QE, (as we defined them) have very different data complexity, e.g., for first-order logic, query evaluation is in AC^0, model expansion is NP-complete.
Database query ϕ can be viewed as a *relation definable with respect to a structure* (a database). See, e.g., the classic textbook by Enderton [4]. The action of defining such a relation is always deterministic – it defines one relation. This relation is a set of tuples, it is not a part of the vocabulary.

3.3 Model Expansion as a Binary Relation

When we talk about powerful solvers (such as those that allow us to compute, say, 3-Colourability or to come up with solutions to logistics problems), we usually produce multiple solutions. In data graphs, we may have multiple children of the same data node. Later in this paper, we will talk about a transition system with nondeterministic actions and a calculus of binary relations. Model expansion (as in Definition 5, the Find part) gives us such a binary relations on structures.

EXAMPLE 5. 3-Colourability is a binary relation such that $(\mathcal{A}, \mathcal{B})$ is in this relation if and only if $\mathcal{A}$ contains an interpretation of the edge and vertex relations (that is, a particular graph), and $\mathcal{B}$ contains an interpretation of relational symbols R, B, G that represents a particular correct colouring of this graph. Note that both $\mathcal{A}$ and $\mathcal{B}$ may also interpret other relational symbols, but those

interpretations do not matter for this relation. Those extra things may be called garbage of the computation.

3.4 The Four Tasks in Applications

In database research, QE has already been studied extensively, e.g. for first-order logic (Codd's relational algebra) and its fragments (e.g. for conjunctive queries), as well as for DATALOG and its variants. SAT has demonstrated its practical importance mostly for propositional logic, a logic of a very low expressive power. Indeed, the success in SAT solving (achieved both by smart algorithms and an exponential growth in hardware capabilities) is one of the most remarkable achievements of logic in Computer Science. Due to this success, the complexity class NP is often called "the new tractable". However beyond the propositional case, the great majority of logics that are interesting and useful in Computer Science are undecidable. For instance, first-order logic is undecidable even in the finite (by the Trakhtenbrot's theorem). In addition, integration of theories often presents as a major problem in Knowledge Representation and in Satisfiability Modulo Theory because a combination of two theories is often undecidable. However, in practice, a finite domain is often given on the input, and, in such a case, the undecidability problem for combinations of theories does not arise.

Moreover, systems for logics with a very high complexity of satisfiability, often perform very well in practice. The explanation is that those systems solve MX, not SAT, since an (active) domain is a part of the input. While SAT continues to be important for propositional logic, the importance of this task for expressive logics used in practice is *greatly overrated*. Just as the query evaluation problem is prevalent in database research, model expansion is very common in the general area of constraint solving. Most constraint solving paradigms solve MX as the main task, e.g. logistics, supply chain management, etc. Java programs, if they are of input-output type, can be viewed as model expansion tasks, regardless of what they do internally. Most combinatorial problems are of that form, many software programs (e.g., web services) and hardware devices (e.g., circuits) are of that form. The Logistics Service Provider in Example 2 has, e.g., customer requests as an input, and routes and packing solutions as outputs. In Example 1, one can have e.g. edges of a graph on the input to formula (2), and colours on the output. ASP systems, e.g., Clasp [6] mostly solve model expansion, and so do CP languages such as Essence [5], as shown in [10]. Problems solved in ASP competitions are mostly in model expansion form. CSP in the traditional AI form (respectively, in the homomorphism form) is representable by model expansion where mappings to domain elements (respectively, homomorphism functions) are expansion functions.

In the algebra we present next, we can view each algebraic expressions as a *network of inter-connected solvers and databases*, jointly solving one task: satisfy all the components. We deal with two types of objects:

- Static Objects: model checking modules M_p:
  - collections of databases
  - *decision* procedures of any kind and complexity, e.g. is a given graph 3-colourable?
  - relations or collections of objects
- Dynamic Objects: model expansion modules M_a:
  - actions, changes
  - combinatorially complex *search* and *optimization* problems, e.g. planning, scheduling, TSP
  - any binary relations on (sets of) structures
  - roles in Description Logics
  - data links in graph databases
  - causality links

4. ALGEBRA WITH INFORMATION FLOW

We now describe a transformation from the "flat" algebra to the "dynamic" one, which gives rise to a modal logic. All we do is we **add information propagation**. Some atomic modules serve as propositions. They are unchanged by the transformation. Other become actions. In each atomic action-module M, we underline designated input symbols and denote them σ_M. Output symbols are those that are free (are not quantified) in the algebraic expression where M occurs. They are denoted ε_M. Thus, **we force multi-dimensional expressions into binary**. For compound expressions α, we use σ_α to denote the union of all input symbols that occur free, and σ_ε to denote all output symbols of α that occur free.

EXAMPLE 6. Consider again the HC-2Col example:

$$\exists Y[M_{\mathrm{HC}}(\underline{V},\underline{X},Y) \wedge M_{2\mathrm{Col}}(\underline{V},Y,Z,T)]. \qquad (3)$$

The quantified symbol Y is not visible from the outside. The output vocabulary of this compound modular system is $\varepsilon = \{Z,T\}$ (for the two colours), the input vocabulary is $\sigma = \{V,X\}$ (for the vertices and edges). In general, *any* direction of information propagation can be specified. For the external (free) symbols, a particular specification of inputs and outputs determines which problem we are solving. For the internal symbols (those that are quantified), it does not matter which symbols are inputs, which ones are outputs. For instance, the internal symbol Y can be considered as an input to the second module, or it can be a symbol on the output who's value is guessed and checked to satisfy both modules.

4.1 Minimal Syntax of the Dynamic Algebra

Fix a vocabulary of atomic module symbols τ_M. Let τ_P, where $\tau_P \subseteq \tau_M$, be a vocabulary of atomic module symbols where inputs are *not* specified.

We call them *propositions*. Alternatively, we can think of these modules as having outputs that are identical to the inputs. Let τ_{act}, where $\tau_{act} \subseteq \tau_M$, be a vocabulary of atomic module symbols $M_i(\underline{X_{i_1}},\ldots,X_{i_k})$, where inputs are underlined. We call them *actions*. For one module symbol M_i, we can potentially have both a proposition, e.g. M_i, and several actions, depending on the choice of the inputs.

We define a calculus of binary relations as follows.

$$\alpha ::= \bot | M_i? | M_a | Z_j | \alpha \cup \alpha | -\alpha | \pi_\delta \alpha | \sigma_\Theta \alpha | \mu Z_j.\alpha \quad (4)$$

Here, M_i are propositions, M_a are actions. Notice that the operations are exactly like in the first algebra.
Variables Z_j range over actions. As usual, we require that Z_j occurs positively (under an even number of $-$) in $\mu Z_j.\alpha$. Requirements on $\pi_\delta \alpha$ and $\sigma_\Theta \alpha$ are as in the "flat" algebra.

4.2 Semantics of the Dynamic Algebra

We interpret the calculus of binary relations over graphs with data points that are relational databases, or graph databases, or any other data structures representable using structures. Our data graphs are transition systems.

DEFINITION 7 (TRANSITION SYSTEM $\mathcal{T}$). *Transition system*

$$\mathcal{T} := (V; (M_a^{\mathcal{T}})_i, (M_p^{\mathcal{T}})_j)$$

(parameterized by valuation $(\mathcal{V}, v)$ defined above) has domain V which is the set of all τ-structures over a fixed domain, which is given by $\mathcal{V}$, and it interprets all actions M_a as subsets of $V \times V$ denoted by $M_a^{\mathcal{T}}$, and all monadic propositions M_p by structures (now nodes in the transition graph) $M_p^{\mathcal{T}} \subseteq V$.

Module variables Z_j that occur free in α are interpreted as actions, i.e., as subsets of $V \times V$.
As before, we require that for any two τ-structures $\mathcal{A}_1$, $\mathcal{A}_2$ which coincide on $vvoc(M_i)$, we have $\mathcal{A}_1 \in \mathcal{V}(v, M_i)$ iff $\mathcal{A}_2 \in \mathcal{V}(v, M_i)$.
Recall how valuations $(\mathcal{V}, v)$ work. First, v maps relational variables to symbols of the vocabulary τ. Second, $\mathcal{V}$, parameterized with v, provides an interpretation to each atomic module, which is a set of structures as before (i.e., a concrete module). In particular, $\mathcal{V}$ also specifies a domain.
We define the extension $[\![\alpha]\!]^{\mathcal{T},\mathcal{V},v}$ of formula α in $\mathcal{T}$ under valuation $(\mathcal{V}, v)$ inductively as follows.

$[\![\bot]\!]^{\mathcal{T},\mathcal{V},v} := \varnothing$.
$[\![M_i?]\!]^{\mathcal{T},\mathcal{V},v} := \{(\mathcal{B},\mathcal{B}) \in V^{\mathcal{T}} \times V^{\mathcal{T}} \mid \mathcal{B} \in \mathcal{V}(v, M_i)\}$.
$[\![M_a]\!]^{\mathcal{T},\mathcal{V},v} := \{(\mathcal{B}_1,\mathcal{B}_2) \in V^{\mathcal{T}} \times V^{\mathcal{T}} \mid \mathcal{B}_1|_{\tau \setminus \varepsilon_{M_a}} = \mathcal{B}_2|_{\tau \setminus \varepsilon_{M_a}}$ and $\mathcal{B}_2 \in \mathcal{V}(v, M_a)\}$.
$[\![\alpha_1 \cup \alpha_2]\!]^{\mathcal{T},\mathcal{V},v} := [\![\alpha_1]\!]^{\mathcal{T},\mathcal{V},v} \cup [\![\alpha_2]\!]^{\mathcal{T},\mathcal{V},v}$.
$[\![-\alpha_2]\!]^{\mathcal{T},\mathcal{V},v} := V^{\mathcal{T}} \times V^{\mathcal{T}} \setminus [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v}$.
$[\![\mu Z_j.\alpha]\!]^{\mathcal{T},\mathcal{V},v} := \bigcap \{R \subseteq V^{\mathcal{T}} \times V^{\mathcal{T}} : [\![\alpha]\!]^{\mathcal{T},\mathcal{V}[Z:=R],v} \subseteq R\}$.
$[\![\pi_\delta(\alpha)]\!]^{\mathcal{T},\mathcal{V},v} := \{(\mathcal{B}_1,\mathcal{B}_2) \in V^{\mathcal{T}} \times V^{\mathcal{T}} \mid$
$\exists \mathcal{C}_1 \exists \mathcal{C}_2\ ((\mathcal{C}_1,\mathcal{C}_2) \in [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v}, \mathcal{C}_1|_\delta = \mathcal{B}_1|_\delta$ and $\mathcal{C}_2|_\delta = \mathcal{B}_2|_\delta)\}$.

$[\![\sigma_{L_1 \equiv L_2}(\alpha)]\!]^{\mathcal{T},\mathcal{V},v} := \{(\mathcal{B}_1,\mathcal{B}_2) \in V^{\mathcal{T}} \times V^{\mathcal{T}} \mid$ 3 cases:
1) $(\mathcal{B}_1,\mathcal{B}_2) \in [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v}$ and $\{L_1, L_2\} \subseteq \sigma_\alpha$ and $\mathcal{B}_1 \models L_1 \equiv L_2$
2) $(\mathcal{B}_1,\mathcal{B}_2) \in [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v}$ and $\{L_1, L_2\} \subseteq \varepsilon_\alpha$ and $\mathcal{B}_2 \models L_1 \equiv L_2$
3) $L_1 \in \sigma_\alpha$ and $L_2 \in \varepsilon_\alpha$ and
$\exists \mathcal{C}\ (\ (\mathcal{C},\mathcal{B}_2) \in [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v}$ and $\mathcal{B}_1|_{\tau \setminus \{L_1\}} = \mathcal{C}|_{\tau \setminus \{L_1\}}$ and
$(\mathcal{B}_1 \models L_1$ iff $\mathcal{B}_2 \models L_2)\ )$.

Here, $\models$ is the standard satisfaction relation as in the first-order logic. Case 3 expresses feedback from output L_2 to input L_1, similar to a feedback in boolean circuits, also used in [11]. Notice that L_1 is a new guessed symbol, so the number of models in the third case may increase. Cases 1 and 2 potentially reduce the number of models.

4.2.1 Satisfaction Relation for the Calculus of Binary Relations

DEFINITION 8 (SATISFACTION, DYNAMIC ALGEBRA). *Given a well-formed algebraic expression α defined by (4), we say that transition system $\mathcal{T}$ and pair of states $(\mathcal{A},\mathcal{B})$* satisfy *$\alpha$ under valuation $(\mathcal{V}, v)$, notation*

$$\mathcal{T}, (\mathcal{A},\mathcal{B}) \models_{(\mathcal{V},v)} \alpha,$$

if $(\mathcal{A},\mathcal{B}) \in [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v}$.

4.2.2 Atomic Modules-Actions

Here, we clarify the semantics of the atomic modules-actions. According to the semantics, atomic actions produce a replica of the current database except the interpretation of the expansion (output) vocabulary changes as specified by the action. This is similar to the inertia law in the Situation Calculus and other formalisms for reasoning about actions.

EXAMPLE 7. To illustrate transitions using our examples, in (3), first $M_{\mathrm{HC}}(\underline{V}, \underline{X}, Y)$ makes transition by producing possibly several Hamiltonian Circuits. The interpretation of the output vocabulary, $\{Y\}$ changes, everything else is transferred by inertia. Then each resulting structure is taken as an input to 2-Colouring, $M_{2\mathrm{Col}}(\underline{V}, Y, Z, T)$, where edges in the cycle, Y, are "fed" to the second argument of $M_{2\mathrm{Col}}$, although this is hidden from the outside observer by the existential quantifier in (3). The second module produces non-deterministic transitions, one for each generated colouring.

4.3 Useful Operations

We introduce some definable operations.

4.3.1 Basic set-theoretic operations and equivalence

$$\begin{aligned}
&\top := -\bot,\\
&\alpha_1 \cap \alpha_2 := -(-\alpha_1 \cup -\alpha_2),\\
&\alpha_1 - \alpha_2 := -(-\alpha_1 \cup \alpha_2),\\
&\alpha_1 \equiv \alpha_2 := (-\alpha_1 \cup \alpha_2) \cap (-\alpha_2 \cup \alpha_1).
\end{aligned}$$

4.3.2 Projection onto the inputs

$$\downarrow \alpha \ := \pi_{\sigma_\alpha} \alpha.$$

This operation is also called "projection onto the first element of the binary relation". It identifies the states in V where there is an outgoing α-transition. Thus,

$$[\![\downarrow \alpha]\!]^{\mathcal{T},\mathcal{V},v} = \{(\mathcal{B},\mathcal{B}) \in V^{\mathcal{T}} \times V^{\mathcal{T}} | \exists \mathcal{B}' \ \ (\mathcal{B},\mathcal{B}') \in [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v}\}.$$

4.3.3 Projection onto the outputs

$$\uparrow \alpha \ := \pi_{\varepsilon_\alpha} \alpha.$$

4.3.4 Unary negation

$$\sim \alpha \ :=\downarrow (-\alpha).$$

It says "there is no outgoing α-transition". By this definition,

$$[\![\sim \alpha]\!]^{\mathcal{T},\mathcal{V},v} = \{(\mathcal{B},\mathcal{B}) \in V^{\mathcal{T}}{\times}V^{\mathcal{T}} \mid \forall \mathcal{B}' \ (\mathcal{B},\mathcal{B}') \notin [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v}\}.$$

By these two definitions, $\downarrow \alpha \ = \sim\sim \alpha$.

4.3.5 Diagonal

$$D :=\sim \bot.$$

$$[\![D]\!]^{\mathcal{T},\mathcal{V},v} = \{(\mathcal{B},\mathcal{B}) \in V^{\mathcal{T}}{\times}V^{\mathcal{T}}\}.$$

This operation is sometimes called the "*nil*" action, or it can be seen as an empty word which is denoted ε in the formal language theory.

4.3.6 Sequential composition

It is very common, in modal logics of programs (e.g. Dynamic Logic), in expressive Description Logics, in graph databases, etc., to consider the composition operator, but not intersection. Sequential composition ($\alpha_1 \circ \alpha_2$) is not definable using the other operations, but is a particular case of intersection ($\cap$), and it can be obtained by imposing a simple sufficient syntactic restriction on the expressions combined.

$$[\![\alpha_1 \circ \alpha_2]\!]^{\mathcal{T},\mathcal{V},v} := \{(\mathcal{A},\mathcal{B}) \in V^{\mathcal{T}}{\times}V^{\mathcal{T}} \mid \exists \mathcal{C}((\mathcal{A},\mathcal{C}) \in [\![\alpha_1]\!]^{\mathcal{T},\mathcal{V},v} \text{ and } (\mathcal{C},\mathcal{B}) \in [\![\alpha_2]\!]^{\mathcal{T},\mathcal{V},v})\}.$$

If there are neither output interference nor cyclic dependencies, then intersection becomes sequential composition:

PROPOSITION 3. *If* $\varepsilon_{\alpha_1} \cap \sigma_{\alpha_2} \neq \varnothing$, $\varepsilon_{\alpha_2} \cap \sigma_{\alpha_1} \neq \varnothing$ *and* $\varepsilon_{\alpha_1} \neq \varepsilon_{\alpha_2}$, *then*

$$\alpha_1 \cap \alpha_2 \ = \ \alpha_1 \circ \alpha_2.$$

4.3.7 Counting

This operation comes from graph databases. It represents a path that is composed of k pieces α, where $n \leqslant k \leqslant m$ and $n < m$.

$$\alpha^{n,m} := \underbrace{\alpha \circ \cdots \circ \alpha}_{n} \circ \underbrace{(\alpha \cup D) \circ \cdots \circ (\alpha \cup D)}_{m-n},$$

where we have a composition of n times α and $m-n$ times $\alpha \cup D$. This definition produces:

$$[\![\alpha^{n,m}]\!]^{\mathcal{T},\mathcal{V},v} = \bigcup_{k=n}^{m} ([\![\alpha]\!]^{\mathcal{T},\mathcal{V},v})^k.$$

4.3.8 Reverse

This operation is common in e.g. Description Logics as well as in graph databases. It amounts to changing the direction of information propagation, i.e., flipping inputs and outputs. It is not definable in the syntax as presented here, but notice that the operations of assigning inputs σ_α and outputs ε_α are silently present in the language (we added them to the "flat" algebra). They could be made explicit, and that would give us the reverse operator.

4.3.9 Subexpression Tests

These operations check if a path in the transition graph starts and ends with the same or different data values.

$$\begin{aligned} \alpha_{=} \ &:= \ \downarrow \alpha \ \equiv \uparrow \alpha, \\ \alpha_{\neq} \ &:= \ -(\downarrow \alpha \ \equiv \uparrow \alpha). \end{aligned}$$

By these definitions,

$$\begin{aligned} [\![\alpha_{=}]\!]^{\mathcal{T},\mathcal{V},v} &= \{(\mathcal{B}_1,\mathcal{B}_2) \in V^{\mathcal{T}}{\times}V^{\mathcal{T}} \mid (\mathcal{B}_1,\mathcal{B}_2) \in [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v} \text{ and } \mathcal{B}_1 = \mathcal{B}_2\}, \\ [\![\alpha_{\neq}]\!]^{\mathcal{T},\mathcal{V},v} &= \{(\mathcal{B}_1,\mathcal{B}_2) \in V^{\mathcal{T}}{\times}V^{\mathcal{T}} \mid (\mathcal{B}_1,\mathcal{B}_2) \in [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v} \text{ and } \mathcal{B}_1 \neq \mathcal{B}_2\}. \end{aligned}$$

4.3.10 Constant Tests

For a (constant) relation R on the domain elements, we can check if R is (or is not) the interpretation of a particular relational variable (under a variable assignment $(\mathcal{V}, v)$) using the selection operation.

$$\begin{aligned} R^{=} \ &:= \ \sigma_{X=\text{`}R'} D, \\ R^{\neq} \ &:= \ -(\sigma_{X=\text{`}R'} D). \end{aligned}$$

By these definitions,

$$\begin{aligned} [\![R^{=}]\!]^{\mathcal{T},\mathcal{V},v} &= \{(\mathcal{B},\mathcal{B}) \in V^{\mathcal{T}}{\times}V^{\mathcal{T}} \mid \mathcal{B} \models_{(\mathcal{V},v)} X \equiv \text{`}R'\}, \\ [\![R^{\neq}]\!]^{\mathcal{T},\mathcal{V},v} &= \{(\mathcal{B},\mathcal{B}) \in V^{\mathcal{T}}{\times}V^{\mathcal{T}} \mid \mathcal{B} \not\models_{(\mathcal{V},v)} X \equiv \text{`}R'\}. \end{aligned}$$

4.4 Two-Sorted Syntax, $L\mu\mu$

The grammar (4) for the algebra with information flow can be equivalently represented in a "two-sorted" syntax, denoted $L\mu\mu$, where expressions for state formulae ϕ and processes α are defined by mutual recursion.

$$\begin{aligned} \alpha &::= \bot \mid M_a \mid Z_j \mid \alpha \cup \alpha \mid -\alpha \mid \pi_\delta(\alpha) \mid \sigma_\Theta(\alpha) \mid \phi? \mid \mu Z_j.\alpha \\ \phi &::= M_i \mid X_i \mid \ \phi \vee \phi \mid \neg\phi \mid \langle\alpha\rangle\phi \mid \mu X_j.\phi \end{aligned} \tag{5}$$

Notice that the second line corresponds to the mu-calculus $L\mu$. We define the necessity modality through the possibility modality: $[\alpha]\phi \ := \ \neg\langle\alpha\rangle\neg\phi$. Thus, we can write $\langle\alpha\rangle\phi$ (respectively, $[\alpha]\phi$) to express that after some (respectively, all) executions of modular system α, property ϕ holds. As usual, $\phi_1 \wedge \phi_2 := \neg\phi_1 \vee \neg\phi_2$. Notice that we have binary (for processes) and unary (for state formulae) fixed points.
The formulae in this logic allow one to specify the goals of the execution, both eventual and extended in time. Thus, $L\mu\mu$ can act as a programming language.

4.5 Semantics of $L\mu\mu$

The modal logic $L\mu\mu$ (5) is interpreted over the same transition system as the calculus of binary relations (4).

State Formulae: Exactly like in the μ-calculus:

$$[\![M_i]\!]^{\mathcal{T},\mathcal{V},v} := \{\mathcal{A} \in V^{\mathcal{T}} \mid \mathcal{A} \in \mathcal{V}(v, M_i)\},$$
$$[\![\phi_1 \vee \phi_2]\!]^{\mathcal{T},\mathcal{V},v} := [\![\phi_1]\!]^{\mathcal{T},\mathcal{V},v} \cup [\![\phi_2]\!]^{\mathcal{T},\mathcal{V},v},$$
$$[\![\neg\phi]\!]^{\mathcal{T},\mathcal{V},v} := \{\mathcal{A} \in V^{\mathcal{T}} \mid \mathcal{A} \notin [\![\phi]\!]^{\mathcal{T},\mathcal{V},v}\},$$
$$[\![\langle\alpha\rangle\phi]\!]^{\mathcal{T},\mathcal{V},v} := \{\mathcal{A} \in V^{\mathcal{T}} \mid \exists\mathcal{B}\ ((\mathcal{A},\mathcal{B}) \in V^{\mathcal{T}}\times V^{\mathcal{T}} \text{ and } (\mathcal{A},\mathcal{B}) \in [\![\alpha]\!]^{\mathcal{T},\mathcal{V},v} \text{ and } \mathcal{B} \in [\![\phi]\!]^{\mathcal{T},\mathcal{V},v})\},$$
$$[\![\mu Z_j.\phi]\!]^{\mathcal{T},\mathcal{V},v} := \bigcap\{R \subseteq V^{\mathcal{T}} : [\![\phi]\!]^{\mathcal{T},\mathcal{V}[Z:=R],v} \subseteq R\}.$$

Process Formulae: Exactly like in the one-sorted syntax, and, in addition, like in Dynamic Logic:

$$[\![\phi?]\!]^{\mathcal{T},\mathcal{V},v} := \{(\mathcal{A},\mathcal{A}) \in V^{\mathcal{T}}\times V^{\mathcal{T}} \mid \mathcal{A} \in [\![\phi]\!]^{\mathcal{T},\mathcal{V},v}\}.$$

4.5.1 Example: Equality Test

EXAMPLE 8 (EQUALITY TEST). Formula $\langle\alpha_1 \equiv \alpha_2\rangle\top$ specifies the set of states from where some executions of α_1 and α_2 lead to the same data value. Here, $\top$ is an abbreviation for any tautology, e.g. $M \vee \neg M$. By the semantics, the meaning of this formula is:

$$[\![\langle\alpha_1 \equiv \alpha_2\rangle\top]\!]^{\mathcal{T},\mathcal{V},v} = \{(\mathcal{B},\mathcal{B}) \in V^{\mathcal{T}}\times V^{\mathcal{T}} \mid \exists\mathcal{B}'\exists\mathcal{B}''\ (\ (\mathcal{B},\mathcal{B}') \in [\![\alpha_1]\!]^{\mathcal{T},\mathcal{V},v} \text{ and } (\mathcal{B},\mathcal{B}'') \in [\![\alpha_2]\!]^{\mathcal{T},\mathcal{V},v} \text{ and } \mathcal{B} = \mathcal{B}')\}.$$

This operation corresponds to an operation used in graph databases, XPath. Non-equality test has a negation ($-$) in front of the equivalence ($\equiv$).

4.5.2 Satisfaction Relation for $L\mu\mu$

DEFINITION 9 (SATISFACTION RELATION, $L\mu\mu$). *Given a well-formed state formula ϕ and process formula α as defined by (5), we say that transition system $\mathcal{T}$ and state $\mathcal{A}$* satisfy *ϕ under valuation $(\mathcal{V}, v)$, notation*

$$\mathcal{T}, \mathcal{A} \models_{(\mathcal{V},v)} \phi,$$

if $\mathcal{A} \in [\![\phi]\!]^{\mathcal{T},\mathcal{V},v}$. For process formulae α, the definition of the satisfaction relation is exactly as in Definition 8.

4.5.3 Two-Sorted = Minimal Syntax

The two representations of the algebra (one-sorted and two-sorted) are equivalent, as we show below. The statement is analogous to the one in [1].

THEOREM 1. *For every state formula ϕ in two-sorted syntax (5) there is a formula $\hat{\phi}$ in the minimal syntax (4) such that $\mathcal{T}, \mathcal{B} \models_{(\mathcal{V},v)} \phi$ iff $\mathcal{T}, (\mathcal{B},\mathcal{B}) \models_{(\mathcal{V},v)} \downarrow\hat{\phi}$, and for every action formula α there is an equivalent formula $\hat{\alpha}$ in the minimal syntax.*

PROOF. We need to translate all the state formulae into process formulae. We do it by induction on the structure of the formula. Atomic constant modules and module variables remain unchanged by the transformation. However, monadic variables are considered binary.

- If $\phi = \phi_1 \vee \phi_2$, we set $\hat{\phi} := \hat{\phi_1} \cup \hat{\phi_2}$.
- If $\phi = \neg\phi_1$, we set $\hat{\phi} :=\sim \hat{\phi_1}$.
- If $\phi = \langle\alpha_1\rangle\phi_1$, we set $\hat{\phi} := \hat{\alpha_1} \circ \hat{\phi_1}$.
- If $\phi = \mu X.\phi_1$, we set $\hat{\phi} := \mu X. \downarrow \hat{\phi_1}$.

All process formulae α except test $\phi_1?$ remain unchanged under this transformation. For test, we have:

- If $\alpha = \phi_1?$, we set $\hat{\alpha} :=\downarrow \hat{\phi_1}$.

It is easy to see that, under this transformation, the semantic correspondence holds. □

4.6 Some Notable Fragments

Since our logic is very expressive, it is not surprising that many logic are fragments of it. For example, the well-known temporal logics CTL, LTL, CTL* are obviously fragments of $L\mu\mu$ since they are fragments of the mu-calculus $L\mu$. What is interesting, however, is to analyze examples of diverse nature and origin where efficient reasoning is the goal. Often, information propagation plays a role there, and a modal logic is obtained as a result. We already saw examples of operations used in graph databases. Graphs are, obviously, binary relations. We now consider several other well-known logics.

4.6.1 Dynamic and Description Logics

Dynamic Logic was created for reasoning about programs.

PROPOSITION 4. *The Dynamic Logic*

$$\begin{array}{l}\alpha ::= M_a \mid \alpha \cup \alpha \mid \alpha \circ \alpha \mid \phi? \mid \alpha^* \\ \phi ::= M_i \mid \phi \vee \phi \mid \neg\phi \mid \langle\alpha\rangle\phi\end{array} \quad (6)$$

is a fragment of (5).

PROOF. By Proposition 3, sequential composition is a particular case of intersection. The other operations of Dynamic Logic are a subset of those in (5). Thus, it is sufficient to express α^*. We have $\alpha^* := \mu Z.(D \cup Z \circ \alpha)$. □

The following property follows from the well-known connection between expressive Description Logic and Dynamic Logic with reverse operator [3].

COROLLARY 1. *Description Logic $\mathcal{ALCI}_{reg}$*

$$\begin{array}{l}R ::= P \mid R \cup R \mid R \circ R \mid id(C) \mid R^* \mid R^- \\ C ::= A \mid C \vee C \mid \neg C \mid \exists R.C\end{array} \quad (7)$$

is a fragment of (5).

In (7), C denotes concepts, R denotes roles, and A and R stand for atomic concepts and roles, respectively. Notation $id(C)$ stands for test, R^- is reverse operator, $\exists R.C$ is the modal "exists" operator.

4.6.2 From Second-Order to First-Order

Here, we explain how atomic modules become standard predicates in the sense of first-order logic. Propositional logic is a fragment of first-order logic. To see this, view propositions as predicate symbols over one-element domain. Similarly, first-order logic is a fragment of second-order logic. In second-order logic, we allow quantification over relations, i.e., over sets of tuples of domain elements. If every such set is a singleton, and every relation we quantify over is unary, then second-order quantifiers behave as quantifiers over domain elements, and second-order collapses to first-order. Our modules are sets of structures. When second-order logic collapses to first-order, **each structure becomes a tuple of domain**

elements, and each module becomes a relation in the sense of first-order logic. In this case, also, third-order fixed points that represent sets of structures, collapse to the usual fixed points that are relations (sets of tuples). Second-order logic is a great tool that allows us to talk about many things uniformly.

4.6.3 Datalog$^{+}_{-}$ as a Modal Logic

Just as our calculus, Datalog$^{+}_{-}$ [2] forces multi-dimensional expressions into binary. This is not immediately apparent, unless one carefully examines the rules in search of explicit or implicit existential quantifiers. Those quantifiers produce possibility modalities, which turn into necessity under negation. For example, the following program

$$emp(X) \rightarrow \exists Y\, hasMgr(X,Y), emp(Y)$$
$$person(P) \rightarrow \exists F\ fatherOf(F,P)$$
$$fatherOf(F,P) \rightarrow person(F)$$

translates into the following formula in our calculus:

$$(emp(X) \rightarrow \langle hasMgr(\underline{X},Y)\rangle\ emp(Y))$$
$$\wedge(person(P) \rightarrow \langle fatherOf(F,\underline{P})\rangle\ \top)$$
$$\wedge([fatherOf(F,\underline{P})]\ person(F)),$$

where atomic module symbols are elements of the set $\delta := \{emp, hasMgr, person, fatherOf\}$, and implication ($\rightarrow$) is the usual abbreviation.

The semantics of Datalog$^{+}_{-}$ is given by the least model that satisfies the rules (which are in the Horn form). By a well-known construction, this least model is expressible by a simultaneous fixed point $\mu\delta.\phi$. The most common reasoning task for Datalog$^{+}_{-}$ is *certain* reasoning, which amounts to computing if a query is true in every expansion of a given database that satisfies a given Datalog$^{+}_{-}$ program.

The language of Datalog$^{+}_{-}$ can be greatly enriched following $L\mu\mu$ (5), by allowing compound expressions inside the modalities, e.g. regular expressions or the special-purpose operations we gave as examples.

5. QUERIES, MACHINES, MODALITIES

In this section, we connect two declarative ways of specifying problems, as in database query answering and as in temporal logic model checking, with each other and with a machine-based approach that has an imperative flavour. The connection between the three approaches is possible because of information propagation.

5.1 Structures as Computing Devices

5.1.1 Abstract Machines

We now introduce very **simple abstract machines**, similar to automata. The main (and only) operation of these machines is a task broadly solved in practice, the Model Expansion task [9], see Definition 5. Our computing devices are *τ-structures*. All they do is **store information and expand**. Declarative specifications of modular systems (formulae) are the programs for these simple devices, and the modal logic from the previous section can be viewed as a programming language for these machines. Program *execution* consists of *constructing a transition system*.

5.1.2 Transition System Determined by α

Without loss of generality, assume that all atomic modules are represented by binary relations, as in the semantics of the calculus of binary relations (4). To talk about all possible executions of α, we construct a transition system TS_α by starting from $\mathcal{T} := (V; (M_a^{\mathcal{T}})_i, (M_p^{\mathcal{T}})_j)$, as in Definition 7, and adding labelled edges produced by each subformula.

DEFINITION 10 (TRANSITION SYSTEM TS_α). *Given formula α in the calculus of binary relations and valuation $(\mathcal{V}, v)$, where v maps $vvoc(\alpha)$ to τ, the* labelled transition system TS_α that represents possible executions of α *is*

$$\mathsf{TS}_\alpha := (V; L),$$

where V is the set of τ-structures, and L is a set of labelled edges. The edges are constructed according to the following rule: If α_i is a subformula of α, then

$$(\mathcal{B}_1, \mathcal{B}_2) \in L \quad \textit{iff} \quad (\mathcal{B}_1, \mathcal{B}_2) \in [\![\alpha_i]\!]^{\mathcal{T},\mathcal{V},v},$$

and the label of $(\mathcal{B}_1, \mathcal{B}_2)$ is α_i.

Notice that since valuation $(\mathcal{V}, v)$ is given, function $\mathcal{V}$ specifies the domain and interpretations of the atomic symbols. Thus, generating TS_α is a constructive process.

5.1.3 Complexity Measures

The time and space complexity of constructing TS_α is associated with the complexity of satisfying α over a given domain (which is the Model Expansion task). While data, expression and combined complexity are considered the main measures of the amount of computation required, we argue that output complexity (in the sense of output-sensitive algorithms) is hugely important as well because it affects the size of the transition system and the number of steps required. Recall that in the basic labelling algorithm that is in the foundation of symbolic model checking, three parameters are multiplied: the number of vertices, the number of edges, and the size of the formula. Data complexity is responsible for the number of vertices, expression complexity measures the size of the formula, and output complexity is responsible for the number of edges. We believe that *input width* and *output width* of a formula should be considered separately. The former affects data complexity, the latter affects output complexity. If we consider all these parameters in interaction, we may be able explain, e.g., why some algorithms for PSPACE-complete problems (such as model checking) work reasonably well in practice, while some algorithms for provably polynomial time problems behave very badly.

5.1.4 Executions for a Given Input

When a particular input structure $\mathcal{A}$ is given, a concrete execution materializes. In this case, we can connect reachability in TS_α with executing α in the following sense.

DEFINITION 11 (REACHABILITY IN TS_α).

We say that state formula ϕ is reachable from the initial state $\mathcal{A}$ by the execution of α, notation $\mathsf{REACH}_\alpha(\mathcal{A}, \phi)$, if, in the transition system constructed by executing α, there is an edge, labelled with α, from the state $\mathcal{A}$ to a state where ϕ holds.

From the construction of TS_α, it follows that one has to construct the edges for all the subformulae of α, in order to construct the edge for α.
We will be interested in the case where ϕ is a conjunction of ground atoms $\bigwedge \bar{\mathcal{E}}$.

5.2 General Evaluation Problem EV

Recall that model checking is a special case of model expansion where the input structure interprets the entire vocabulary (second-order variable vocabulary in our case). We now define another problem that is essentially model checking – a higher-order counterpart of the Query Evaluation Problem QE.

DEFINITION 12 (Evaluation Problem EV).
Given:
1. *valuation $(\mathcal{V}, v)$,*
2. *formula $\phi(\bar{X})$ in the calculus without information propagation,*
3. *σ-structure $\mathcal{A}$, where $\mathcal{A} = (A; \bar{R}_\sigma)$ and $\bar{R}_\sigma$ interprets a part of the visible (free) relational variables of ϕ,*
4. *a tuple of relations $\bar{\mathcal{E}}$ that interprets the rest of the visible (free) relational variables of ϕ.*

Find: *Structure $\mathcal{B}$ such that*

$$\overbrace{(\underbrace{A, \bar{R}_\sigma}_{\mathcal{A}}, \bar{R}, \bar{\mathcal{E}})}^{\mathcal{B}} \models_{(\mathcal{V},v)} \phi(\bar{X}, \bar{Y})[\ \bar{R}_\sigma/\bar{X},\ \bar{\mathcal{E}}/\bar{Y}\ ]?$$

Tuple of relations $\bar{R}$ interprets the internal (not free) relational variables of ϕ. We require, as usual, that for the substitutions, the arities have to coincide. Note that σ may be empty. The domain of $\mathcal{A}$ is the same as the one given by $\mathcal{V}$, and σ must be a subset of τ, which is the concrete vocabulary provided by v for the relational variables.

Notice that **this task imposes a direction of information propagation**, thus it turns an expression in a "flat" algebra to one in the "dynamic" one. The General Evaluation Problem (EV) is equivalent to the Model Checking Problem (MC) for a formula where the internal relational symbols are second-order-existential quantified. Model Expansion (MX) for $\phi(\bar{X})$ is equivalent to first guessing the output relations $\bar{\mathcal{E}}$ and then using EV to check.

PROPOSITION 5. *Query Evaluation problem QE is a particular case of the General Evaluation problem EV.*

PROOF. The statement follows immediately from our explanation about why first-order logic is a fragment of second-order in Subsection 4.6.2. □

REMARK 4. *The definitions of the General Evaluation task EV and the Model Expansion task MX could be given with partial interpretations of the predicates on the input (e.g. in terms of sets of ground atoms or 3-valued structures). That version is more convenient in other contexts, but is not necessary here.*

5.3 Temporal Logic Tasks: temp-MC, temp-SAT

Here we introduce counterparts of Model Checking and Satisfiability tasks in the context of modal temporal logics.

DEFINITION 13 (MODEL CHECKING: temp-MC).
Given: valuation $(\mathcal{V}, v)$, transition system $\mathcal{T}$, σ-structure $\mathcal{A}$, where $\sigma \subseteq \tau$, state formula ϕ. Decide: $\mathcal{T}, \mathcal{A} \models_{(\mathcal{V},v)} \phi$.

A common version of this problem is the one where one is asked to compute all the states where the formula is true. This version is used in practical model checking algorithms for temporal logics. We restrict our attention to the problem of finding some structure of that sort.

DEFINITION 14 (temp-MC-SEARCH). *Given: valuation $(\mathcal{V}, v)$, transition system $\mathcal{T}$, state formula ϕ. Find: structure $\mathcal{A}$ such that*

$$\mathcal{T}, \mathcal{A} \models_{(\mathcal{V},v)} \phi.$$

The transition system on the input is determined by the valuation and, in a way, is redundant. The valuation also gives the domain of $\mathcal{A}$.
And here is another important problem that looks similar on the surface, but can be drastically different computationally.

DEFINITION 15 (SATISFIABILITY: temp-SAT).
Given: State formula ϕ, valuation v that fixes a concrete vocabulary. Find: valuation $\mathcal{V}$, transition system $\mathcal{T}$, structure $\mathcal{A}$ (one of the states of $\mathcal{T}$), such that

$$\mathcal{T}, \mathcal{A} \models_{(\mathcal{V},v)} \phi.$$

The main difference here is that we need to find a valuation $\mathcal{V}$, which determines the domain and the interpretation of "unary" and "binary" module symbols. The transition system $\mathcal{T}$ is then constructed from those.

5.4 Connecting Machines, Queries, Modalities

In this subsection, we show that, in our translation from "flat" to modal logic, the temp-MC task is the same as the expansion-evaluation task EV, and is also equivalent to the reachability in the execution graph. Surprisingly, assigning input and outputs to the internal variables does not matter.

THEOREM 2. *Suppose we are given, on the input, a formula ϕ in the "flat" algebra, $(\mathcal{V}, v)$, $\mathcal{A}$ and $\bar{\mathcal{E}}$. For any assignments of inputs and outputs to the internal (not free) variables $\bar{R}$ of ϕ that produces α from ϕ, we have:*

$$\overbrace{\mathcal{T}, \mathcal{A} \models \langle\alpha\rangle \bigwedge \bar{\mathcal{E}}}^{\text{temp-MC}} \Leftrightarrow \mathsf{REACH}(\mathcal{A}, \bigwedge \bar{\mathcal{E}}) \Leftrightarrow \overbrace{(\mathcal{A}, \bar{R}, \bar{\mathcal{E}}) \models \phi(\bar{\mathcal{E}})}^{\text{EV}}$$

It follows that temporal logic model checking (symbolic, SAT-based, etc.) can be used for solving the Evaluation problem. Vise versa, techniques developed for query evaluation, may influence model checking algorithms, for the corresponding fragments.

PROOF. We explain the main idea. The rest follows from the definitions of the tasks. Suppose a direction of information propagation for the internal variables has been assigned. The transition system TS_α is constructed bottom up on the structure of the formula. We first guess the extensions of the input variables and identify the set of states with those guessed extensions. Then we propagate the information according to each of the modules-actions, from inputs to the outputs. This step gives us the extensions of all the output variables of the atomic modules, and the corresponding states. The construction proceeds up on the structure of the formula, adding more transitions. This process is equivalent to guessing all the internal variables first, then checking if the guess satisfies the atomic modules, then proceeding as before. Thus, the order of information propagation for the internal variables does not matter. □

We can think of ϕ as a specification of an algorithm, and of a corresponding α, where the information propagation to the internal variable is fixed, as an implementation of it. The direction of information propagation is implementation-dependent. While it is not important in theory, it does affect the practical complexity of algorithms, as it influences the complexity measures discussed in Subsubsection 5.1.3.

A connection between database query evaluation and temporal model-checking has been noticed before. In footnote 10 of the paper "From Church and Prior to PSL" [14], Moshe Vardi wrote that the two problems are "analogous", and that "the study of the complexity of database query evaluation started about the same time as that of model checking". However, we are not aware of any papers where a precise correspondence has been established. This correspondence is closely related to an important question we are going to discuss next.

5.5 On Robust Decidability of Modal Logics

Modal logic, starting from the simplest modal logic ML where the necessity and possibility modalities are added to the propositional logic, to much more complex logics with path and state quantifiers and fixed points such as CTL, LTL, $L\mu$, are known for their good computational properties. Moshe Vardi posed the question to identify the main reasons for the robust decidability of modal logics, and partially answered it [13]. In a paper with the same title [7], Erich Grädel discussed the problem further. He states the main motivation for this research as: "We would like to have more powerful logics than ML, CTL and even the μ-calculus that retain the nice properties of modal logics." Vardi also wrote that "... modal logic, in spite of its apparent propositional syntax, is essentially a first-order logic, since the necessity and possibility modalities quantify over the set of possible words ..." model-checking problem for the modal logic ML can be solved in linear time, while satisfiability is PSPACE-complete. On the other hand, validity, and thus, satisfiability, is highly undecidable.

Previous partial explanations of robust decidability of modal logics are related to two-variable fragment of first-order logic FO^2, finite model property, tree model property, guarded fragments of first-order fragments, bisimulation-invariance and the characterization theorems. However, all of these explanations view modal logics as interpreted over Tarskian structures. Then states are viewed as domain elements of a "flat" Tarskian structure, propositions are predicates ranging over domain elements. Still, in all these explanations, the complexity of satisfiability remains hugely different in the modal and in the classical setting.

While this subject deserves a separate paper, we provide a brief explanation here. In our understanding, possible words are structures, and the modalities quantify over those, not over domain elements.

We believe that the following is needed for a computationally good behaviour of a logic:

- information flow,
- an initial structure.

The initial structure can just be a domain. Of course, if we have a domain, the task is no longer Satisfiability, but Model Expansion. For the same logic, the latter task is of significantly lower complexity. In the propositional case, we get the domain for free, as is explained below, even when we consider Satisfiability task.

We would now like to draw an analogy with physics. The process of producing a modal logic from a classical one is very similar to the well-understood physical process of crystal formation from a liquid matter. The conditions needed for crystals to form requires the presence of a force (of information flow here), and a seed (a structure in our case). During crystal formation, energy gets released. Released energy is also responsible for a special form of perpetual motion in the fascinating concept of time crystals invented by a Nobel laureate Frank Wilczek in 2012.

5.5.1 Propositional Case

By the propositional case, we mean a modal logic built on top of propositional. Typical examples include CTL, LTL, $L\mu$, PDL, but we also allow all the operations of $L\mu\mu$. As in the more general case, the states in the transition system $\mathcal{T}$ are structures.

If a propositional $L\mu\mu$ formula is satisfiable, it is satisfiable over a transition system built from structures over one-element domain. To see this, take an arbitrary domain. We construct equivalence classes of domain elements induced by the subsets of the set of propositional variables. Each such class is fully determined by a structure with one element. Thus, one domain element is enough.

Since we always have a fixed domain, temporal satisfiability is equivalent to the search version of temporal model checking, temp-SAT= temp-MC-SEARCH,

Classical RA	Lifted RA
Basic units are Relations = sets of tuples of domain elements	Basic units are Modules = classes (sets, if the domain is fixed) of structures
Object variables	Relational variables
Query Evaluation task, QE check if $\bar{a}$ is in the relation defined w.r.t. $\mathcal{A}$	Evaluation task, EV check if $\bar{\mathcal{E}}$ is in an expansion of $\mathcal{A}$
Query Computation: mismatch inputs are structures, outputs are relations	Model Expansion, MX: no mismatch both inputs and outputs are structures
hard to connect to modal logics	connection to modal logics is straightforward
Fixed points = sets of tuples, $\mu Z.\phi(Z)$ is a relation	Fixed points = sets of structures, $\mu Z.E(Z)$ is a set of structures

With Information Propagation

Resulting modal logic is "propositional"	Resulting modal logic is "first-order"
States in TS are one-element structures	States in TS are structures

Figure 1: Comparison of classical and lifted setting for Relational Algebra (RA) with recursion

which, in turn, is related to model expansion in the classical logic setting. Thus, the existence of a fixed domain, together with the force of information propagation, explains robust decidability of modal logics that are built on top of propositional ones. With our explanation, all the previous explanations remain true. For instance, we do have only unary and binary predicates, but in a generalized sense. Our guards are generalized too, and so is the tree model property.

6. CONCLUSION

The relational data model has recently been heavily criticized for being outdated, not being able to link data, to the degree that it's been called a "legacy technology". We have shown that those claims are perhaps premature. The same data model that comes from Codd's relational algebra, with two modification that do not affect the main language, can represent relational data, graph databases and operations on them, powerful inter-connected solvers, ontologies, all in the same framework.
While logic is already responsible for great technological advances, we strongly believe that the next technological shift can only happen if we change the units we operate on, as we proposed here.

7. REFERENCES

[1] Faried Abu Zaid, Erich Grädel, and Stephan Jaax. Bisimulation safe fixed point logic. In Rajeev Goré, Barteld P. Kooi, and Agi Kurucz, editors, *Advances in Modal Logic 10, invited and contributed papers from the tenth conference on "Advances in Modal Logic," held in Groningen, The Netherlands, August 5-8, 2014*, pages 1–15. College Publications, 2014.

[2] Andrea Calì, Georg Gottlob, Thomas Lukasiewicz, Bruno Marnette, and Andreas Pieris. Datalog+/-: A family of logical knowledge representation and query languages for new applications. In *Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, 11-14 July 2010, Edinburgh, United Kingdom*, pages 228–242. IEEE Computer Society, 2010.

[3] Diego Calvanese and Giuseppe De Giacomo. Expressive description logics. In Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors, *The Description Logic Handbook: Theory, Implementation, and Applications*, pages 178–218. Cambridge University Press, 2003.

[4] Herbert B. Enderton. *A mathematical introduction to logic*. Academic Press, 1972.

[5] Alan M. Frisch, Warwick Harvey, Chris Jefferson, Bernadette Martínez-Hernández, and Ian Miguel. Essence: A constraint language for specifying combinatorial problems. *Constraints*, 13:268–306, 2008.

[6] Martin Gebser, Benjamin Kaufmann, and Torsten Schaub. Conflict-driven answer set solving: From theory to practice. *Artificial Intelligence*, 187-188:52–89, August 2012.

[7] Erich Grädel. Why are modal logics so robustly decidable? In *Current Trends in Theoretical Computer Science*, pages 393–408. 2001.

[8] Antonina Kolokolova, Yongmei Liu, David Mitchell, and Eugenia Ternovska. On the complexity of model expansion. In *Proc., 17th Int'l Conf. on Logic for Programming, Artificial Intelligence and Reasoning (LPAR-17)*, pages 447–458. Springer, 2010. LNCS 6397.

[9] D. G. Mitchell and E. Ternovska. A framework for representing and solving NP search problems. In *Proc. AAAI'05*, pages 430–435, 2005.

[10] D. G. Mitchell and E. Ternovska. Expressiveness and abstraction in ESSENCE. *Constraints*, 13(2):343–384, 2008.

[11] S. Tasharrofi and E. Ternovska. A semantic account for modularity in multi-language modelling of search problems. In *Proceedings of*

the 8th International Symposium on Frontiers of Combining Systems (FroCoS), pages 259–274, October 2011.

[12] M. Y. Vardi. The complexity of relational query language. In *14th ACM Symp.on Theory of Computing, Springer Verlag (Heidelberg, FRG and NewYork NY, USA)-Verlag*, 1982.

[13] Moshe Y. Vardi. Why is modal logic so robustly decidable? In Neil Immerman and Phokion G. Kolaitis, editors, *Descriptive Complexity and Finite Models, Proceedings of a DIMACS Workshop 1996, Princeton, New Jersey, USA, January 14-17, 1996*, volume 31 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 149–184. DIMACS/AMS, 1996.

[14] Moshe Y. Vardi. From Church and Prior to PSL. In Orna Grumberg and Helmut Veith, editors, *25 Years of Model Checking - History, Achievements, Perspectives*, volume 5000 of *Lecture Notes in Computer Science*, pages 150–171. Springer, 2008.